

A bootstrap method for estimating bias and variance in statistical multispecies models using highly disparate data sets

Lorna Taylor, Verena M. Trenkel, Vojtech Kupca, Gunnar Stefansson¹

Abstract

Statistical multispecies models of multiarea marine ecosystems use a variety of data sources to estimate parameters using composite or weighted likelihood functions with associated weighting issues and questions on how to obtain variance estimates. Regardless of the method used to obtain point estimates, a method is needed for variance estimation. A bootstrap technique is introduced for the evaluation of uncertainty in such models, taking into account inherent spatial and temporal correlations in the data sets thus avoiding many model–specification issues, which are commonly transferred as assumptions from a likelihood estimation procedure into Hessian–based variance estimation procedures. The technique is demonstrated on a real data set and used to look for estimation bias and the effects of different aggregation levels in population dynamics models.

Key words: bootstrapping, fish population dynamics, nonlinear models, correlated data

¹Author for correspondence, Science Institute, University of Iceland, Taeknigardur, Dunhaga 5, 107 Reykjavik, Iceland, gstefans@gmail.com

1 Introduction

One of the tasks of a statistical model is to consolidate data from various sources by using them simultaneously to estimate parameters. Within an ecosystem context this is one of the functions of the Gadget modelling environment, as originally proposed by Stefansson and Palsson (1998), described in Begley (2004) and subsequently used in multiple applications (Björnsson and Sigurdsson, 2003; Taylor et al., 2007). The importance of using all data in a single model has been emphasised by several authors (Demyanov et al., 2006; Methot, 1989) and although the benefits of using all available data in a single model-fitting procedure are clear, it is certainly not without problems, including the question of variance estimation and weighting of data sources.

Variance estimation is important when estimating parameters of any statistical model. This becomes critical when the interest is in obtaining confidence statements from complex models of the population dynamics of exploited marine species. In this setting, multiple data sources with widely different properties are routinely used in the estimation process. This paper demonstrates a novel use of bootstrapping to address such issues. The approach is generic, but it has special application to statistical models of (multiple and interacting) marine populations such as those developed within the Gadget framework. The protocol to estimate likelihood component weights and optimise model parameters is given in Taylor et al. (2007) and the weighting protocol is based on that described in Stefansson (1998) & Stefansson (2003)

Variances of parameters in nonlinear models have traditionally been derived from an inverted Hessian matrix at the optimum, when the method of least squares (or maximum likelihood) is used for estimation. Specifically, when a single data source is used and a sum of squares (SS) is the appropriate fitting criterion, it is well known that the Hessian (of the SS at the optimum) or Jacobian matrices (of the residuals) can be used to obtain estimates of uncertainty, i.e. variance estimates for the parameters obtained from the minimisation.

For statistical inference such as confidence statements to hold in this scenario several conditions need to be satisfied, simply put: the model needs to be correct. As for linear models, an assumption of normality is required for Hessian-based inference, although this is not necessary for point estimation. Variance assumptions (i.e. homoscedasticity and knowledge of the ratios of variances in individual data sets) are also important.

Naturally, alternative likelihoods can in principle be used and this approach has been extensively developed in the theory of generalized linear models (McCullagh and Nelder, 1989). However, when the distributional properties of the data are not well understood or the models are incorrect, these approaches will fail as seen in several examples in fishery science (Patterson et al., 2001).

When multiple data sources are considered, a different issue arises, which is how the data sets should be weighted. Given the problems associated in using single data sources, it would seem rather futile to try to use the analytic or parametric approaches mentioned above for uncertainty estimation with multiple data sources (Stefansson, 1998). Point estimates can of course still be obtained using any popular method, such as ordinary least squares (OLS) or some variant of reweighted least squares. Notably, Stefansson (2003) suggests a reweighting scheme designed to detect or accommodate model misspecification in the context of the objective function and this is used by Taylor et al. (2007).

An alternative approach to variance estimation when using simple (single, homogeneous) data sets is to use a non-model-based method such as the bootstrap (Efron, 1979; Efron and Tibshirani, 1994). Bootstrap methods provide a general nonparametric mechanism for estimating uncertainty in any estimation method. In introductory textbooks it is assumed that the data are simple measurements without correlation. However, semi-parametric approaches have also been developed to sample residuals from a model, possibly from a distribution (parametric bootstrap) and even with a correlation structure (Davison and Hinkley, 1997).

Methods of estimating variances in fish stock assessment models have been evaluated by many authors including Gavaris et al. (2000) and Patterson et al. (2001). The fact that the bootstrap methods appear to do well in these comparisons should indicate they are a promising avenue to investigate, but the multivariate nature of the data under consideration here needs to be taken into account. The importance of the correlation structure has been noted earlier and disregarding it has been found to potentially lead to entirely incorrect conclusions in single-species assessments (Myers and Cadigan, 1995).

A typical fisheries data base consists of data from a variety of sources. Every sample from each data source can be classified according to sampling location and time. A model such as Gadget operates on certain time-steps and spatial units which are always fairly large but with more than one time-step per year and commonly more than one area. Within any modelled spatio-temporal unit there will normally be several data samples.

If the data with which the model is parameterised are to be bootstrapped, an immediate concern is therefore what the sampling unit should be for any resampling method. A unit of measurement in marine studies tends to be based on a single fish and elementary resampling might bootstrap on individual fish (as in e.g. Gudmundsdóttir et al., 1988). This pretty much assumes that all individually measured fish are independent which is invalid for several reasons, the simplest of which is that fish samples have an internal correlation structure reflecting the biology. Fish of similar size and age tend to be found in similar locations so the correlation structure of e.g. length distributions is not reflected by the multinomial distribution corresponding to simple random sampling (Hrafnkelsson and Stefansson, 2004). Resampling entire samples of fish can potentially be used to account for the fact that samples of ages and lengths at the same station (a survey tow) will be correlated through the intra-haul correlation (Pennington and Volstad, 1994). This is not quite enough, however, since fish at close locations will also tend to be similar due to a fine-scale spatial structure which can not be easily modelled (e.g. Stefansson and Palsson, 1997).

These correlations are exacerbated when multiple samples from commercial vessels are considered, since a sequence of samples may be from a single school or cluster of schools of similar fish.

In addition to the above problem, one needs to take into account the variety of data sources. Biological samples from commercial catches may be collected on a fine temporal and spatial scale whereas scientific surveys are typically only conducted once or twice a year and may or may not completely overlap spatially. Other data sets such as species composition of stomach contents or tagging experiments may be collected at completely different resolutions to age or length data.

The following sections describe a methodology to store and handle data in such a manner as to permit bootstrapping for the purpose of variance estimation.

This bootstrapping methodology is subsequently used to evaluate the effect of variable aggregation levels in statistical multispecies models. This is important since simple issues such as the choice of the number of size classes, areas or time steps has a profound effect on computational time in multispecies models. It is seen that these concerns can be quite easily addressed using the proposed methodology which includes a specific approach to database design and bootstrapping from the database.

2 Data handling in multispecies modelling

When modelling population dynamics it should be recognised that different (spatial and/or temporal) scales may be required to answer different questions. To allow for spatial structural flexibility, data storage for modelling purposes is therefore best implemented on an aggregation level somewhat finer than needed for most modelling purposes. There is no good reason to go to the level of individual fish for modelling population dynamics. On the contrary, there is every reason to store data aggregated to spatial and temporal cells where between-cell correlations can be assumed to be negligible. These aggregated cells are then referred to as elementary data units.

In this paper simple choices are made in the low-level aggregation in order to at least reduce these correlations. For example, entire length samples are only used combined (rather than lengths of individual fish), then size-related correlations within samples are avoided (Hrafnkelsson and Stefansson, 2004). Similarly, data in the data base are always aggregated within fairly large areas and the shortest time-step is one month. This approach should eliminate effect of the intra-haul correlation (Pennington and Volstad, 1994) and those correlations between age-groups (Myers and Cadigan, 1995) which are related to local behaviour.

Once data are available in a data base form, where all data have been pre-summarised to the required spatial and temporal cells (the data units), any model definition corresponding to an aggregation of these data units also corresponds to simple aggregates of the data units. Thus one obtains a simple function which transforms a collection of data elements into data for modelling. Rerunning the model based on a different aggregation scheme (e.g. different areas, temporal scales or length groupings) becomes a fairly trivial task since the extraction routines from the data base to the Gadget input file formats are automated.

The aggregation method from the standardised data base to Gadget input files varies somewhat depending on the data source. Some data, e.g. length distributions, are simply added whereas others consist of e.g. mean length at age and go through a computational mechanism.

A suggestion for a standardised fisheries data base, which can be used in this manner is described in Kupca and Sandbeck (2003) and Kupca (2004a).

Figures 1-2 describe the spatial definition of data units (known as subdivisions) around Iceland and adjacent waters developed by Taylor (2003). The spatial structure is based mainly on bathymetry, hydrography and species assemblages with some further disaggregation defined by fishing regulations.

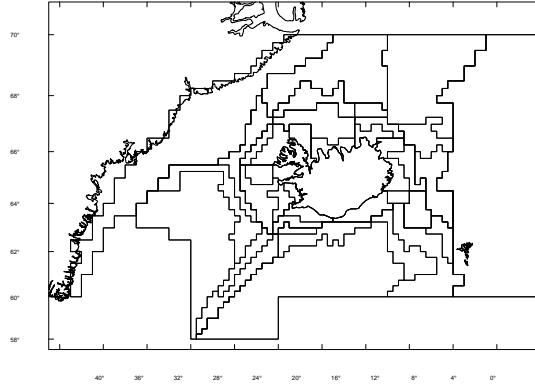


Figure 1: The spatial structure of data storage for Icelandic and adjacent waters. These areas are referred to as ‘subdivisions’.

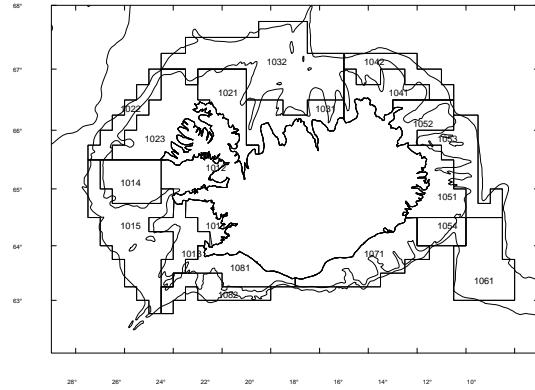


Figure 2: The Icelandic shelf subdivisions plotted along with 200m and 500m depth contours.

3 Bootstrapping from a data base

Naturally, any one of the aggregation methods described above can be based on any collection of data elements. In particular, a model aggregation can be based on resampling (with replacement) entire data cells from the data base. Each such resampling will lead to a new model data set.

It would therefore seem appropriate to base the resampling directly on data which have already been blocked into “data units” in the standardised data base. It is thus proposed that resampling be implemented by allowing sampling with replacement from the data base.

A typical model run for parameter estimation based on such a data set will result in a resampled parameter estimate. The collection of all such estimates form a bootstrap sample.

Details of an implementation of such bootstrap extraction utilities from a data base are described in Kupca (2004b). It should be noted that such a procedure can not be trivially implemented as a standard data base access procedure, as repeated extractions of the same elements of a table are required for bootstrapping but disallowed by standard SQL statements.

With this approach the “sampling unit” in the context of resampling becomes the entire set of all measurements within a spatio–temporal “data unit” in the data base. This will not be a simple n –dimensional vector of measurements but a “ragged array” of widely different data types, and commonly not all data types will appear in a given bootstrap sample.

With this approach the “sampling unit” in the context of resampling becomes the entire set of measurements within a spatio–temporal data unit in the data base. This will not be a simple n –dimensional vector of measurements but a “ragged array” of widely different data types, and commonly not all data types will appear in a given bootstrap sample.

With this approach the “sampling unit” in the context of resampling becomes the entire set of data units within a spatio–temporal model area in the data base. This will not be a simple n –dimensional vector of measurements but a “ragged array” of widely different data types, and commonly not all data types will appear in a given bootstrap sample.

An approach similar to the one proposed, but with much simpler data and simpler models was developed in Hannesson et al. (2008) where both simulated and real data are used to verify that this general approach is likely to work for the simple case of variation only coming in through tagging data.

4 A fisheries example

4.1 The setting

The example marine system used in this paper is based on cod in Icelandic waters (fig. 2) with an approach very similar to Taylor et al. (2007). Two surveys are used to monitor the stock, in spring and fall, giving population indices as well as biological samples. Landings information is available from official data bases and raw biological data (length distributions, age compositions) along with survey data in MRI data bases (see e.g. MRI, 2004; Palsson et al., 1989; Sigurdsson et al., 1997; Taylor et al., 2007, for a description of data and surveys).

Routine assessments of this stock tend to use a time and age scale of a year and aggregate data into a single area (MRI, 2004). These assessments do not explicitly model body growth but directly include measurements of e.g. mean weight at age. The present model includes a body growth model which is implemented by “moving” fish from a source length group into several other length groups so as to obtain the appropriate average growth. In this setting a scale needs to be chosen for modelling fish length. If this scale is too coarse in relation to the time scale then on each time-step some fish will not grow whereas others will inevitably grow too much, resulting in an incorrect distribution of length at age for older age groups.

As pointed out by Vandermeer (1978), there is a balance to be found between estimation errors due to too small size classes and distribution error caused by too large size classes. It is therefore of particular interest to investigate the effects of the choice of scale on outputs and this can be done with the approach proposed here.

4.2 The data set

The model is a parametric and deterministic forward simulation model. A single simulation results in a complete population structure, including predictions of all data sets, as described in Begley (2004) and Taylor et al. (2007) and a corresponding evaluation of a (negative log-)likelihood function (sums of squares in the present paper).

With the exception of landings data, data sets are only used in the likelihood components. For simplicity, landings data are used directly in the population models, whereby the populations are simply reduced in numbers to be in accordance with the corresponding landed weight.

The likelihood data used are:

- Length distributions from the March survey (1985 — 2003), October survey (1995 — 2003) and sampling from the commercial fishery (1984 — 2003).
- Age-length frequencies from the March survey (1989 — 2003), October survey (1995 — 2003) and sampling from the commercial fishery (1984 — 2003).
- Survey indices from the March survey (1985 — 2003) and October survey (1995 — 2003). These are calculated from the length distributions and are disaggregated (“sliced”) into three groups which correspond roughly to age 1, age 2 and age3+
- The ratio of immature:mature by length group from the March survey (1985 — 2003).

4.3 Population models, likelihoods and parameters

When population abundance is estimated, one particularly important facet of likelihood functions is how survey abundance numbers (U) are compared to the modelled population numbers (N). The simplest approach is to assume a multiplicative relationship viz $E[U] = qN$ where the unknown q is commonly termed the catchability coefficient. In many cases it is seen that this is inadequate due to non-Gaussian behaviour as well as apparent nonlinearity in the relationship (e.g. Stefansson (1992)). On the other hand, estimation of parameters in $E[U] = qN^\beta$ or (more appropriately) $\ln(U) = \alpha + \beta \ln(N) + \text{error}$ can result in poorly determined power/slope coefficients (Taylor et al., 2007).

Simple confidence intervals for the slope commonly imply them to be (significantly) different from 1.0 but without a fixed pattern (Taylor et al., 2007, and references therein) and it follows that the choice of which slopes to estimate and which to fix at a value of 1.0 is more an art than a science. A common approach, however, is to fix the slope to 1.0 at older ages and to estimate it for the youngest ages. This takes into account possible nonlinearities due to behavioural differences as well as possible density dependent natural mortality, both of which are (more) likely for the youngest ages (e.g. Myers and Cadigan, 1993).

The models are:

Model 1 All survey indices with the power fixed to 1. The mean length of recruits and mean length at age of the initial population defined.

Model 2 The power for the survey index estimated for length groups 1 and 2 but fixed to 1 for length group 3. The mean length of recruits and mean length at age of the initial population defined.

Model 3 All survey indices with the power fixed to 1. The mean length of recruits and mean length at age of the initial population calculated from the growth parameter estimated in the optimisation procedure.

The parameters estimated are:

- one parameter for the growth function
- one parameter for the length update (i.e. growth transition matrix)
- one parameter for each fleet selection pattern (commercial catch, March survey and October survey)
- two parameters defining the maturation ogive (which are correlated)
- the number of recruits (abundance of age 1) for each year (1984 — 2003)
- the abundance at ages 2 — 11 at the start of the model

4.4 Estimation protocol

The weights on the likelihood components are calculated for each model (i.e. each bootstrap run), according to the protocol described in Taylor et al. (2007).

For Model 2, an initial optimisation run is done to move the parameters into an appropriate range before the weights are estimated. Otherwise the starting parameters are arbitrary.

Basically, the bootstrapping approach consists of the following.

- The base data are stored in a standardised data base:
 - Time aggregation: 1 month
 - Spatial aggregation: subdivision
 - Further disaggregation is based on a range of categories including fishing gear, fishing vessel class, sampling type (e.g. harbour, sea and survey).
- To bootstrap the data, the list of subdivisions is sampled (with replacement) and stored.
- The list of resampled subdivisions then used to extract data (with replacement so the same data set may be repeated several times in a given bootstrap sample).
- For a single bootstrap Gadget model, the same list of resampled subdivisions is used to extract each likelihood dataset i.e. length distributions, survey indices and age-length frequencies are extracted from the same spatial definition.
- The full dataset is extracted and 100 bootstrap datasets.

It should be emphasised that when the resampling estimates are obtained, the entire estimation procedure is repeated for each bootstrap sample. Since the estimation procedure includes an iterative reweighting scheme, this reweighting is repeated for every bootstrap sample. The procedure as a whole is quite computationally intensive but can easily be run in parallel, e.g. on a computer cluster.

In contrast to this approach, Hessian-based approaches usually omit the reweighting of likelihood components when estimating uncertainty. It is by now well-established that incorrect model assumptions can seriously affect parameter estimates through incorrect weights (Stefansson, 2003) and simple Hessian-based methods will not capture this as a part of the uncertainty estimation procedure.

The bootstrap estimation is implemented for the three model structures and for each model structure three levels of time-step aggregation (1 month, 2 month and 3 month blocks) are tested.

Each model structure and time-step aggregation uses equivalent bootstrap data, i.e. for a given time scale, all model structures use the same bootstrapped data.

4.5 Model output

Given the optimised parameter estimates it is possible to output a wide range of descriptors of the model ecosystem as Gadget operates on and stores the number in each age-length cell for each time-step of the model. For this study, the estimated parameters along with derived biomass trajectories (end of year total biomass, age 4+ biomass and spawning stock biomass) are considered.

5 Results

5.1 Single model results

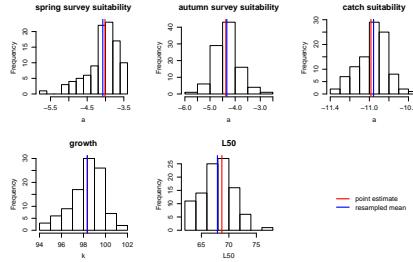


Figure 3: Histograms of some optimised parameters from 100 bootstrap model runs with the point estimate and bootstrap mean indicated.

The simplest (and least interesting) model outputs are the point estimates of model parameters. Fig. 3 gives histograms of bootstrap estimates of several parameters. For each parameter, the point estimate from the full data set and the mean of the bootstrap estimates are also indicated. The differences between the point estimate and the bootstrap mean can be seen to be relatively minor, i.e. there is no obvious sign of an estimation bias. It should be noted that the maturation parameters are correlated, affecting the relationship between the point estimate and bootstrap mean for L50. These plots indicate the results of the model run on monthly time-steps with the survey power estimated for length group 3 (i.e. Model 2) but the results are similar for all other model structures and time-step lengths.

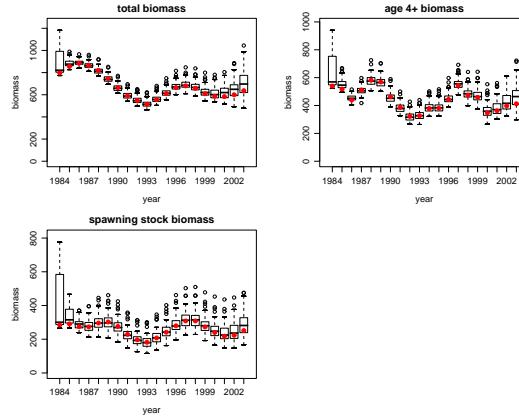


Figure 4: Boxplots of the end of year biomass estimated by the 100 bootstrap models with the point estimate indicated by the red point. The box indicates the interquartile range and the whiskers to the data point which is no more than 1.5 times the interquartile range from the box. Any further outlying data points are indicated as points.

Boxplots can be used (fig. 4) to illustrate bootstrapped trajectories of various abundance or biomass measures. It is seen that the main variation appears in the initial and final years. The initial and final years are of course considerably different from the intermediate ones, but in different ways. The number of fish in the initial year are part of the estimation procedure and therefore of a different nature when compared to subsequent years. Further, the survey starts in 1985 (with the model starting in 1984), which makes the initial conditions somewhat poorly determined. The final years are on the other hand poorly determined since there is relatively little information in the objective function for the younger year classes as they have only been surveyed for a few years.

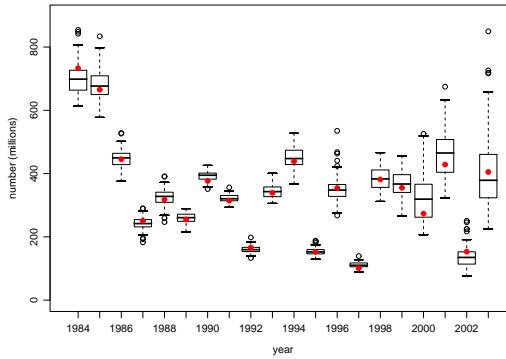


Figure 5: Boxplots of the number of recruits (age 1) in each year estimated by the 100 bootstrap models with the point estimate indicated by the red point.

The same effects are seen for estimated recruitment at age 1 (fig. 5) where there is less variation in the intermediate years than the earliest or later years.

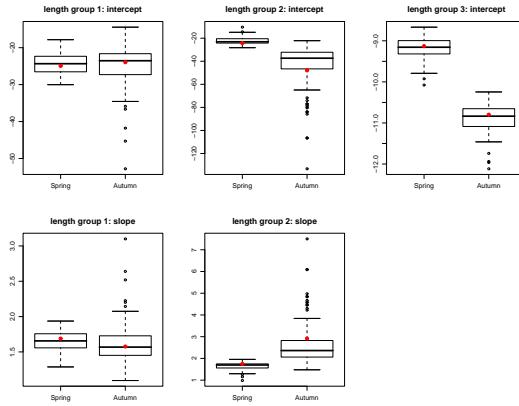


Figure 6: Boxplots of the survey intercept ($\ln(\text{catchability})$) and slope (power) estimated by the 100 bootstrap models for Model 2 with the point estimate indicated by the red point.

When comparing the slope and log-intercept of the observation model across surveys (fig. 6), the difference in basic catchability levels is not of interest as in both cases these are indices, subject to arbitrary scaling. More interesting, however, is a comparison of the variation in the estimates.

There is considerably more variation in the parameter estimates obtained for the autumn survey, which is in accordance with the reduced number of tows in fall (Palsson et al., 1989; Sigurdsson et al., 1997).

5.2 Comparing models with different time-steps

Different time scales may require different length scales. In particular, a coarse time scale of one year may imply that a coarse 5cm length scale may be sufficient but such a length scale will not be adequate for a short time scale. In this particular case, the 1-month (12 step) time scale was used with 1 cm length groups (as in the raw data) but for 3-month (4 step) and 2-month (6 step) models 2 cm groupings were used. Tests were also conducted with 1 cm length groups in the 4- and 6-time-step cases (not shown), giving the same results as the 2 cm groupings presented. The models used to illustrate this comparison are of type Model 1.

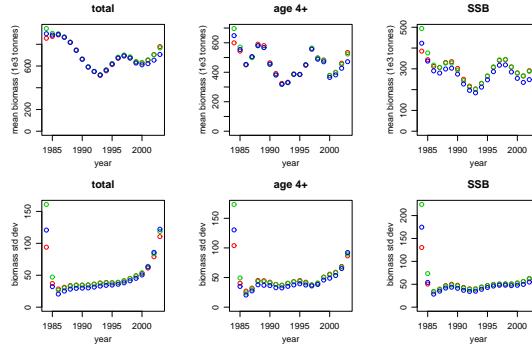


Figure 7: Mean biomass (upper panels) and standard deviation of biomass (lower panels) from models with three different time-step intervals (12, 6 and 4 steps per year).

A comparison of biomass trends (Fig. 7, upper panels) shows little difference in mean predicted biomass or its standard deviation (Fig. 7, lower panels) across the different time scales. The standard deviation is slightly lower for the monthly model (Fig. 7, lower panels). Otherwise differences reflect the uncertainty of the initial year.

Biomass standard deviation (Fig. 7, lower panels) also shows the greater uncertainty for the initial year and increasing uncertainty towards the end of the model run where there is less information on year classes.

These simple tests of scales imply which aggregation levels can be used and these tests are simple to conduct if the process of moving from the data base to the Gadget input data is automated. Naturally, this process has to be automated anyway for the bootstrap procedure proposed in this paper.

5.3 Comparing models with observation model slope fixed or estimated.

The bootstrap methodology can also be used to compare models with different assumptions, in this case to investigate the effect of using variable vs fixed slopes in the relationship between the survey index and model abundance. All model structures are considered (c.f. Section 4.3) and the effect of fixing the mean length at age of the initial population and recruits, rather than calculating

the length using the estimated growth parameters, is also evaluated. While the figures illustrate results from models with 4 time-steps per year the results are equivalent to those from other time scales.

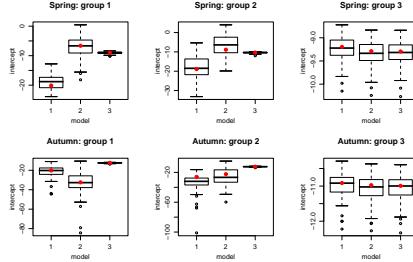


Figure 8: Estimated intercept of the survey index for each length group for the three different models. The red dot indicates the point estimate. Note the difference in scales.

It is seen in fig. 8 that there is a substantial difference in the absolute magnitude of the intercept depending on the assumptions made regarding the mean length at age of the initial population, which is fixed in Model 1 and estimated in Models 2 and 3.

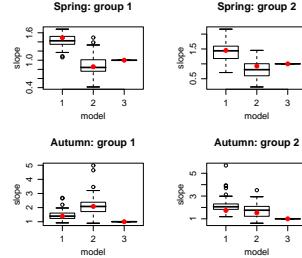


Figure 9: Estimated slope of the survey index for each length group for the three different models. The red dot indicates the point estimate. NB the slope for length group 3 is fixed to 1.

The fact that the difference in intercept magnitude among models is not reflected in the biomasses (below) suggests that the change must be absorbed in the estimated slope as is confirmed by the reduction in the intercept variability when the slope is fixed at 1.0. This is indeed seen in fig. 9 where one notes that higher slopes correspond to lower intercepts from the previous figure.

Naturally, placing different assumptions on the initial year makes a difference in the starting point. It is interesting, however, that these considerably different model assumptions result in largely the same overall stock trajectories apart from the final years as seen in fig. 10.

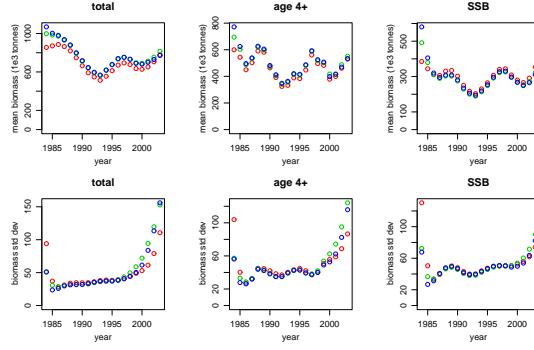


Figure 10: Mean biomass (upper panels) and standard deviation of biomass estimates (lower panels) from models with four time-steps per year.

6 Discussion

Data bases described in this paper enable resampling-based extraction methods. Thus the data bases generate bootstrapped data sets providing uncertainty estimates from highly complex models of marine ecosystems. Since traditional Hessian-based methods have obvious and serious issues due to the assumptions required, the bootstrap approach clearly has considerable advantages (e.g Demyanov et al. 2006).

Several modifications and alternatives to the original bootstrap methodology (Efron, 1979; Efron and Tibshirani, 1994) have been presented. For example, to account for correlations in simple non-replacement sampling schemes (as used for most questionnaires or “sample surveys”), without-replacement bootstraps and with-replacement bootstraps have been suggested along with somewhat more general resampling procedures for complex survey data (McCarthy and Snowden, 1985; Gross, 1980; Rao and Wu, 1988; Sitter, 1992). Theoretical assumptions and derivations behind these approaches do not easily extend to the present situation with disparate data sets, composite likelihoods in the estimation phase and last but not least the highly nonlinear population dynamics models used as a basis for obtaining predicted values and error sums of squares or likelihood functions. The “trick” in the current proposal is not a theoretical development but the methodology of having the bootstrap sampling unit y_i be a “ragged array” (or data “list”) defined in a sufficiently aggregated manner that these objects can be assumed to be independent.

Some of these modifications of the original bootstrap have been developed for marine surveys (Smith, 1997) but this has been intended to reflect e.g. the sampling design used for the surveys and simple estimation of quantities such as a stratified mean. In the present setting the data need to go through an aggregation procedure to be used in a nonlinear population dynamics model and it is the output of this model which is of interest, not variances in the input. Thus there is a need for the bootstrap to mimic this aggregation procedure for the full data from raw data or finer-scale aggregates. This is the case with any population dynamics or assessment model, used in fisheries or other areas of resource harvesting particularly in a multispecies and multi-area context.

With the approach presented here scale changes are a fairly trivial change when data and models are set up as described here, but such changes tend to constitute a major upheaval in traditional modelling environments, where data from different sources are manually adapted to input formats for the models.

A comparison of the effect of using different time-steps does not seem to indicate a considerable effect on biomass trends or corresponding uncertainty estimates except for the initial year. An interesting conclusion from that comparison is that models of this cod stock do not need very fine time scales even if they include growth.

From a technical point of view, the ability to do such scale comparisons without manually modifying input data files should not be underestimated. The difference in using trimesters in place of monthly time-steps is considerable in terms of computation time which becomes an important factor when models start incorporating multiple species on multiple areas and the same applies for the length intervals used since these affect consumption calculations.

In all cases, the results from these bootstrap models indicate that there is no bias in model parameterisation as there is little difference (and no pattern) between the point estimate and bootstrap mean. Even with different model assumptions (e.g. fixing or estimating the power of the survey–abundance relationship) there is remarkably little difference in the modelled biomass trends. Such stability indicates that the population dynamics are well captured by these models and data.

It is reassuring that the modelled years in which the greatest uncertainty is found are those where it is expected i.e. the initial year and then increasing towards the end of the modelled time period. The first year is the most data poor with no survey data or age–length compositions and towards the end of the time period there are fewer cohorts with data available for most ages.

The methodology proposed here is certainly quite computationally intensive. However this is also the case for many other methods. For example, the Hessian–based approach which has been extensively used in fishery science and general ecology needs to be quite carefully implemented since the Hessian matrix needs to be estimated quite accurately (see e.g. Tinker et al., 2006, who use expensive central differencing to estimate the Hessian matrix). For these particular case studies, bootstrapping verified the suitability of less computationally intensive models.

Acknowledgements

Much of the work described here was undertaken while authors LT, VK and GS were employed at the MRI (Marine Research Institute, Reykjavik) and uses data from the MRI databases. The Gadget code has been in development for more than a decade by many programmers at the MRI and IMR (Institute of Marine Research, Bergen). The work was supported in part by EU grants QLK5-CT1999-01609 and FP6 TP8.1 502482, a grant from The Icelandic Centre for Research (Rannis).

References

Begley, J. (2004). Gadget User Manual. Technical Report 120, Marine Research Institute.

Björnsson, H. and T. Sigurdsson (2003). Assessment of golden redfish (*Sebastes marinus* L.) in Icelandic waters. *Scientia Marina* 67(Suppl 1), 301–314.

Davison, A. and D. Hinkley (1997). *Bootstrap Methods and Their Application*. Cambridge University Press.

Demyanov, V., S. N. Wood, and T. J. Kedwards (2006). Improving ecological impact assessment by statistical data synthesis using process-based models. *Journal of the Royal Statistical Society Series C* 55(1), 41–62.

Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics* 7(1), 1–26.

Efron, B. and R. J. Tibshirani (1994). *An Introduction to the Bootstrap*. Chapman & Hall/CRC.

Gavaris, S., K. R. Patterson, C. D. Darby, P. Lewy, B. Mesnil, A. E. Punt, R. M. Cook, L. T. Kell, C. M. O'Brien, V. R. Restrepo, D. W. Skagen, and G. Stefansson (2000). Comparison of uncertainty estimates in the short term using real data. ICES (Int. Counc. Explor. Sea) Document C.M. 2000/V:03.

Gross, S. (1980). Median estimation in sample surveys. *Proceedings of the Section on Survey Research Methods*, 181–184.

Gudmundsdóttir, Á., B. Æ. Steinarsson, and G. Stefánsson (1988). A simulation procedure to evaluate the efficiency of some otolith and length sampling schemes. *ICES (Int. Counc. Explor. Sea) C.M. 1988/D:14* 500, 14.

Hannesson, S., A. Jakobsdóttir, J. Begley, L. Taylor, and G. Stefansson (2008). On the use of tagging data in statistical multispecies multi-area models of marine populations. *ICES Journal of Marine Science*. submitted.

Hrafnkelsson, B. and G. Stefansson (2004). A model for categorical length data from groundfish surveys. *Canadian Journal of Fisheries and Aquatic Sciences* 61(7), 1135–1142.

Kupca, V. (2004a). A standardized database for fisheries data. ICES (Int. Counc. Explor. Sea) Document C.M. 2004/FF:15.

Kupca, V. (2004b). Bootstrapping of areas. In *Final Report: dst²: Development of structurally detailed statistically testable models of marine populations*, Volume 118 of *Technical Report*, pp. xx–xx. Marine Research Institute, Reykjavik.

Kupca, V. and P. Sandbeck (2003). dst² Datawarehouse structure and data import. In *dst²: Development of structurally detailed statistically testable models of marine populations*, Volume 98 of *Technical Report*, pp. 37–46. Marine Research Institute, Reykjavik.

McCarthy, P. J. and C. B. Snowden (1985). The bootstrap and finite population sampling. *Vital Health Stat* 2 95, 1–23.

McCullagh, P. and J. A. Nelder (1989). *Generalized Linear Models* (2 ed.). London: Chapman & Hall/CRC. 511 p.

Methot, R. D. (1989). Synthetic estimates of historical abundance and mortality for northern anchovy. *Amer. Fish. Soc. Symp.* 6, 66–82.

MRI (2004). State of marine stocks in Icelandic waters 2003/2004. Prospects for the quota year 2004/2005. Technical Report 102, Marine Research Institute.

Myers, R. A. and N. G. Cadigan (1993). Density-dependent juvenile mortality in marine demersal fish. *Canadian journal of fisheries and aquatic sciences (Print)* 50(8), 1576–1590.

Myers, R. A. and N. G. Cadigan (1995). Statistical analysis of catch-at-age data with correlated errors. *Canadian journal of fisheries and aquatic sciences (Print)* 52(6), 1265–1273.

Palsson, O. K., S. A. Schopka, G. Stefansson, and B. A. Steinarsson (1989). Icelandic groundfish survey data used to improve precision in stock assessments. *J. Northw. Atl. Fish. Sci.* 9, 53–72.

Patterson, K., R. Cook, C. Darby, S. Gavaris, L. Kell, P. Lewy, B. Mesnil, A. Punt, V. Restrepo, D. W. Skagen, and G. Stefansson (2001). Estimating uncertainty in fish stock assessment and forecasting. *Fish and Fisheries* 2(2), 125–157.

Pennington, M. and J. H. Volstad (1994). Assessing the effect of intra-haul correlation and variable density on estimates of population characteristics from marine surveys. *Biometrics* 50, 1–8.

Rao, J. N. K. and C. F. J. Wu (1988). Resampling inference with complex survey data. *Journal of the American Statistical Association* 83(401), 231–241.

Sigurdsson, T., E. Hjorleifsson, H. Bjornsson, and O. K. Palsson (1997). Fall groundfish survey in icelandic waters (stofnmaeling botnfiska a Islandsmidum haustid 1996). Technical Report 61, Marine Research Institute. (in Icelandic).

Sitter, R. R. (1992). A resampling procedure for complex survey data. *Journal of the American Statistical Association* 87(419), 755–65.

Smith, S. J. (1997). Bootstrap confidence limits for groundfish trawl survey estimates of mean abundance. *Canadian Journal of Fisheries and Aquatic Sciences* 54(3), 616–630.

Stefansson, G. (1992). Notes on the stock-dynamics and assessments of the Icelandic cod. ICES (Int. Counc. Explor. Sea) Document C.M. 1992/G:71.

Stefansson, G. (1998). Comparing different information sources in a multispecies context. In F. Funk, T. Q. II, J. Heifetz, J. Ianelli, J. Powers, J. Schweigert, P. Sullivan, and C. Zhang (Eds.), *Fishery Stock Assessment Models: Proceedings of the international symposium; Anchorage 1997*, 15th Lowell Wakefield Fisheries Symposium, pp. 741–758.

Stefansson, G. (2003). Issues in multispecies models. *Natural Resource Modeling* 16(4), 415–438.

Stefansson, G. and O. K. Palsson (1997). Statistical evaluation and modelling of the stomach contents of Icelandic cod (*Gadus morhua*). *Canadian Journal of Fisheries and Aquatic Sciences* 54(1), 169–181.

Stefansson, G. and O. K. Palsson (1998). A framework for multispecies modelling of boreal systems. *Reviews in Fish Biology and Fisheries* 8, 101–104.

Taylor, L. (2003). Definition of areas in Icelandic waters. In *dst²: Development of structurally detailed statistically testable models of marine populations*, Volume 98 of *Technical Report*, pp. 222–230. Marine Research Institute, Reykjavik.

Taylor, L., J. Begley, V. Kupca, and G. Stefansson (2007). A simple implementation of the statistical modelling framework Gadget for cod in Icelandic waters. *African Journal of Marine Science* 29(2), 223–245.

Tinker, M. T., D. F. Doak, J. A. Estes, B. B. Hatfield, M. M. Staedler, and J. L. Bodkin (2006). Incorporating diverse data and realistic complexity into demographic estimation procedures for sea otters. *Ecol Appl* 16(6), 2293–312.

Vandermeer, J. (1978). Choosing category size in a stage projection matrix. *Oecologia* 32(1), 79–84.